arXiv:1611.00574v1 [cs.FL] 2 Nov 2016

Compositional Reasoning for Shared-variable Concurrent Programs

Fuyuan Zhang[1], Yongwang Zhao[1], David Sanán[1], Yang Liu[1],
Alwen Tiu[1], Shang-Wei Lin[1], Zhimin Wu[1], and Jun Sun[2]

[1] School of Computer Science and Engineering, Nanyang Technological University, Singapore
[2] Singapore University of Technology and Design

Abstract. Scalable and automatic formal verification for concurrent systems is always demanding, but yet to be developed. In this paper, we propose a verification framework to support automated compositional reasoning for concurrent programs with shared variables. Our framework models concurrent programs as succinct automata and supports the verification of multiple important properties. Safety verification and simulations of succinct automata are parallel compositional, and safety properties of succinct automata are preserved under refinements. Formal verification of finite state succinct automata can be automated. Furthermore, we propose the first automated approach to checking rely-guarantee based simulations between infinite state concurrent programs. We have prototyped our algorithm and applied our tool to the verification of multiple refinements.

1 Introduction

Automatic verification of concurrent programs is a challenging task. Due to interleaving, the state space of a concurrent program could grow exponentially, which makes it infeasible to directly reason about the global state space. A promising way of conquering the state explosion problem is compositional reasoning [21,36,28,27,39], which aims at breaking the global verification problems into small localized problems. Extensive research [18,17,25,24,22,12] have been conducted on developing rely-guarantee based automatic verification techniques for safety properties of concurrent programs. However, to ensure that safety properties of concurrent programs are preserved after compilation, it is also necessary to show that the checked programs are refined correctly. To the best of our knowledge, all existing approaches to checking rely-guarantee based simulations of concurrent programs [30] are manual.

In this paper, we propose a framework of automated compositional reasoning for shared-variable concurrent programs, which supports both safety verification and refinement checking. In our framework, concurrent programs are modelled as succinct automata, which can be viewed as an extension of program graphs [3]. A succinct automaton consists of both component transitions, specifying behaviors of a local program, and environment transitions, which overapproximate behaviors of other programs in the environment. The idea of integrating these two types of transitions is the key to ensure parallel compositionality. The development of our framework proceeds in two directions.

The first direction focuses on parallel compositionalities of safety and simulations of succinct automata, which are very useful in developing compositional proof of global properties. For example, our definition of weak simulations between succinct automata allows compositional reasoning through establishing a local refinement relationship. Let SA_1 (resp. $\widehat{SA_1}$) and SA_2 (resp. $\widehat{SA_2}$) be two succinct automata and $SA_1||SA_2$ (resp. $\widehat{SA_1}||\widehat{SA_2}$) be their parallel composition. Since our notion of weak simulation is compositional, we can prove that $SA_1||SA_2$ weakly refines $\widehat{SA_1}||\widehat{SA_2}$ by proving that SA_1 (resp. SA_2) weakly refines $\widehat{SA_1}$ (resp. $\widehat{SA_2}$). As safety properties of succinct automata are preserved under refinements, parallel compositionalities of safety and simulations allow us to extend safety properties of high level concurrent programs to low level concurrent programs in compositional ways.

The second direction aims at automating our compositional reasoning techniques. Formal verification of succinct automata with finite domains can be fully automated. Both safety properties and simulations can be specified as fixed points of corresponding monotonic functions [14], which can be characterized through inference rules. Calculation of the least or greatest fixed points of related inference rules can be automated using worklist algorithms [34]. Due to space limitation, we omit the presentation of these inference rules and focus on checking simulations of infinite state succinct automata. We have developed an SMT-based approach to checking weak simulations for succinct automata with infinite domain. To the best of our knowledge, we are the first to propose automatic verification of rely-guanrantee based simulations for infinite state concurrent programs. We have prototyped our algorithm in F# and verified multiple refinements in automated manner.

Our contributions are three-folds. First, we propose a new formalism, succinct automata, that facilitates automatic verification of multiple properties of shared-variable concurrent programs. Second, we show compositionality results on safety properties and simulations in our framework. Third, we provide an SMT-based approach to verifying simulations for infinite state succinct automata.

2 Related Work

Extensive research have been conducted on the verification of concurrent programs. Basic approaches to conquering the state explosion problem of concurrent systems include (but not limited to) symbolic model checking [6], partial order reduction [38,35,19], abstraction [20,10,13,4,11], symmetry reduction [26,8,16] and compositional reasoning [21,36,28,27,33]. The formalism of succinct automata is inspired by rely-guarantee style reasoning [28,27]. We mainly discuss related work on the compositional reasoning of properties considered in this paper.

Safety Verification. Succinct automata-based approach to safety verification is closest to thread-modular verification [18], where safety properties are characterized by a set of unsafe states and a global system is safe iff unsafe states are not reachable. In this paper, we focus on invariance properties of succinct automata. Checking strong invariants of succinct automata is dual to verifying whether corresponding sets of unsafe states are reachable. Hence, the approach in [18] can be applied to verify strong invariants of (parallel) succinct automata with finite domains. Work in [25,24,22,12]

combined compositional reasoning with abstraction refinement [9]. Moreover, [22,12] allow local variables of different threads to be correlated, which makes their proof rules complete. Developing abstraction refinement techniques under our framework is our future work.

Simulations. Our work on checking weak simulations is related to previous approaches [5,31,7,30] on compositional reasoning of concurrent programs refinement. In [5,31,7], parallel compositionality is achieved by allowing the environments to have arbitrary behaviors, which is considered too strong in general (as was also pointed out by [30]). Our definition of weak simulations for succinct automata is closely related to and inspired by [30], where a rely-guarantee based simulation, called RGSim, for concurrent programs is proposed. Their compositionality rules for RGSim form the basis of a relational proof method for concurrent programs transformations. Our work differs with theirs mainly in that we aim at developing automatic verification of weak simulations between succinct automata. Also, instead of treating all variables as global variables, we distinguish between local variables and global variables. This greatly reduces the state space of local succinct automata. Compared to [23], which has proposed the first automated proof system for refinement verification of concurrent programs, our approach to refinement checking is more general and is not limited to any specific rules of refinement. Work in [29] proposed an automated refinement checking technique for infinite state CSP programs. Their approach is not developed for shared-variable concurrent programs.

3 Succinct Automata

Succinct automata aim to model both local behaviors of a program and its environment in a unified way, and to provide a convenient way to specify useful properties of programs and to support compositional reasoning over them. We distinguish between global variables and local variables when modeling concurrent programs.

3.1 Syntax and Semantics

Let *Dom* be a finite or infinite (numeric) domain and $V = \{v_1, ..., v_n\}$ be a finite set of variables ranging over *Dom*. An *atomic predicate* over V is of the form $f(v_1, ..., v_n) \sim b$, where $f : \mathit{Dom}^n \rightarrow \mathit{Dom}$ is a function, $\sim \in \{=, <, \leq, >, \geq\}$ and $b \in \mathit{Dom}$. A *predicate* over V is a Boolean combination of atomic predicates over V. We write V' for $\{v'_1, ..., v'_n\}$ that refers to variables in V after transitions. Let $\mathcal{F}(V)$ (resp. $\mathcal{F}(V \cup V')$) denote the set of predicates over V (resp. $V \cup V'$). A *valuation* is a function from variables to a domain. Given a predicate $\psi \in \mathcal{F}(V_1)$ and a valuation $\mathbf{v} : V_2 \rightarrow \mathit{Dom}$, where $V_1 \subseteq V_2$, we write $\psi(\mathbf{v})$ to denote that ψ evaluates to true under the valuation $\mathbf{v}$ and write ψ' to stand for the predicate derived from ψ by substituting all its variables in V_1 with corresponding variables in V'_1.

Definition 1. *A Succinct Automaton is a tuple* $SA = (Q, q_0, V, Init, Inv, Env, \Sigma, Edge)$*, where*

– *Q is a finite set of locations and $q_0 \in Q$ is an initial location.*

- $V = V_G \cup V_L$ *and* V_G *(resp.* V_L*) is a finite set of global (resp. local) variables ranging over Dom, where* $V_G \cap V_L = \emptyset$*.*
- $Init \in \mathcal{F}(V)$ *defines initial values of variables at* q_0*.*
- $Inv : Q \to \mathcal{F}(V)$ *constrains the values of variables at each location.*
- $Env : Q \to \mathcal{F}(V_G \cup V'_G)$ *specifies environment transitions at each location.*
- Σ *is a finite set of action labels which includes the silent action* τ*.*
- $Edge \subseteq Q \times \Sigma \times \mathcal{F}(V \cup V') \times Q$ *is a finite set of edges specifying component transitions.*

For each location $q \in Q$, transitions specified by $Env(q)$ are made by the environment when SA stays at location q. An edge is of the form $e = (q, a, \mu, q')$, where μ defines the transition condition and is of the form $\mu := G(V) \wedge \bigwedge_{v'_i \in V'} v'_i = f_i(V)$, where $G(V)$ is a guard for e and f_i is a function $f_i : \mathit{Dom}^n \to \mathit{Dom}$ for $1 \leq i \leq n$. Actions labels in Σ are used when we check weak simulations of succinct automata. The main purpose of Inv is to overapproximate reachable states at each control location of a concurrent program. This also facilitates the formalization of the compatibility condition on succinct automata (introduced later). A succinct automaton is *closed* if its environment cannot modify its global variables.

The semantics of succinct automata is defined as a labeled transition system. A *state* of a succinct automaton is a pair $s = (q, \mathbf{v})$ of location q and valuation $\mathbf{v} : V \to \mathit{Dom}$. We denote with $S_{SA} = \{(q, \mathbf{v}) |\ q \in Q, \mathbf{v} : V \to \mathit{Dom}\}$ the state space of SA. A state $(q, \mathbf{v})$ is an initial state iff $q = q_0$ and $Init(\mathbf{v})$ holds. We say that a predicate ψ is satisfied on $(q, \mathbf{v})$ iff $\psi(\mathbf{v})$ holds.

Let $\mathbf{v}_1 : V_1 \to \mathit{Dom}$ and $\mathbf{v}_2 : V_2 \to \mathit{Dom}$ be two valuations such that $V_1 \cap V_2 = \emptyset$. We define $\mathbf{v}_1 \oplus \mathbf{v}_2 : V_1 \cup V_2 \to \mathit{Dom}$ by $\mathbf{v}_1 \oplus \mathbf{v}_2(v) = \mathbf{v}_1(v)$ for $v \in V_1$ and $\mathbf{v}_1 \oplus \mathbf{v}_2(v) = \mathbf{v}_2(v)$ for $v \in V_2$. Let $\mathbf{v}_G : V_G \to \mathit{Dom}$ (resp. $\mathbf{v}_L : V_L \to \mathit{Dom}$) be valuations over global (resp. local) variables. In the rest of the paper, we also use $(q, \mathbf{v}_G \oplus \mathbf{v}_L)$ to represent a state for convenience. Given a valuation $\mathbf{v} : V \to \mathit{Dom}$, we define $n(\mathbf{v}) : V' \to \mathit{Dom}$ as $n(\mathbf{v})(x') = \mathbf{v}(x)$.

We define two types of transitions, namely *component transitions* and *environment transitions*, for succinct automata. There is a component transition between two states $(q, \mathbf{v}) \xrightarrow{\sigma} (q', \mathbf{v}')$ iff there exists an edge of the form $e = (q, \sigma, \mu, q') \in Edge$ and $Inv(q)(\mathbf{v}) \wedge \mu(\mathbf{v} \oplus n(\mathbf{v}')) \wedge Inv(q')(\mathbf{v}')$ holds. There is an environment transition between two states $(q, \mathbf{v}) \xrightarrow{\phi} (q', \mathbf{v}')$ iff $q = q'$, $Env(q) = \phi$, $Inv(q)(\mathbf{v}) \wedge \phi(\mathbf{v}_G \oplus n(\mathbf{v}'_G)) \wedge Inv(q)(\mathbf{v}')$ and $\mathbf{v}_L = \mathbf{v}'_L$, where $\mathbf{v} = \mathbf{v}_G \oplus \mathbf{v}_L$ and $\mathbf{v}' = \mathbf{v}'_G \oplus \mathbf{v}'_L$. Notice that in an environment transition, only values of global variables can be modified and values of local variables remain unchanged. We say that a component (resp. an environment) transition $(q, \mathbf{v}) \xrightarrow{\sigma} (q', \mathbf{v}')$ (resp. $(q, \mathbf{v}) \xrightarrow{\phi} (q', \mathbf{v}')$) preserves a predicate ψ if we have that $\psi(\mathbf{v})$ implies $\psi(\mathbf{v}')$.

A *run* of SA is a finite or infinite sequence of alternating environment and component transitions starting from an initial state $(q_0, \mathbf{v}_0)$:

$$(q_0, \mathbf{v}_0) \xrightarrow{\phi_0} (q_0, \mathbf{v}'_0) \xrightarrow{\sigma_1} (q_1, \mathbf{v}_1) \xrightarrow{\phi_1} (q_1, \mathbf{v}'_1) \xrightarrow{\sigma_2} (q_2, \mathbf{v}_2) \cdots$$

We say that a predicate ψ is satisfied on a run iff it is satisfied on all states on that run.

A finite *local path* of SA is a sequence of edges $\pi = e_1, ..., e_n$, where $e_i = (q_i, \sigma_i, \mu_i, q'_i)$, $e_n = (q_n, \sigma_n, \mu_n, q'_n)$ and $q'_i = q_{i+1}$ for $1 \leq i < n$.

P_1:
```
while (true) {
   flag1:=1;
   turn:=2;
   await (flag2=0∨turn=1) {
      Critical Section;
   }
   flag1:=0;
}
```

P_2:
```
while (true) {
   flag2:=1;
   turn:=1;
   await (flag1=0∨turn=2) {
      Critical Section;
   }
   flag2:=0;
}
```

Fig. 1. A Simplified Peterson's Algorithm

We write $(q, \mathbf{v}) \rightarrow^* (q', \mathbf{v}')$ if there exists a finite run of SA, (consisting of zero or more transitions), from $(q, \mathbf{v})$ to $(q', \mathbf{v}')$ and say that $(q', \mathbf{v}')$ is *reachable* from $(q, \mathbf{v})$. The set of reachable states of SA is the set of states reachable from initial states of SA. Regarding environment transitions, we write $(q, \mathbf{v}) \xrightarrow{\phi^*} (q, \mathbf{v}')$ to denote a finite sequence of environment transitions of SA starting from $(q, \mathbf{v})$ to $(q, \mathbf{v}')$, where $Env(q) = \phi$. Similarly, for component transitions, we write $(q, \mathbf{v}) \xrightarrow{\tau^*\sigma\tau^*} (q', \mathbf{v}')$ to mean that SA has first taken a finite number of silent actions τ, followed by a component transition labelled by an action σ, and then made another finite number of silent actions.

Example 1 We model a simplified Peterson's algorithm using succinct automata as an example. The pseudo code in Figure 1 shows a simplified version of Peterson's algorithm with two processes P_1 and P_2.

In Figure 2, we model the above two processes as $SA_1 = (Q_1, q_0, V, Init_1, Inv_1, Env_1, \Sigma_1, Edge_1)$ and $SA_2 = (Q_2, p_0, V, Init_2, Inv_2, Env_2, \Sigma_2, Edge_2)$ respectively, where $V = \{\mathtt{flag}_1, \mathtt{flag}_2, \mathtt{critical}_1, \mathtt{critical}_2, \mathtt{turn}\}$, $\Sigma_1 = \{\tau, c_1\}$ and $\Sigma_2 = \{\tau, c_2\}$. Here, we treat all variables as global variables. The automaton SA_1 (resp. SA_2) starts at location q_0 (resp. p_0), where each variable has an initial value of 0, and has five locations q_0, q_1, q_2, q_3 and q_4 (resp. p_0, p_1, p_2, p_3 and p_4). Invariants for locations are presented in ovals. Component transitions are represented by solid line arrows, together with the action labels and predicates on them. We omitted the predicates specifying the variables whose values remain unchanged in component transitions. Environment transitions are represented by dashed line arrows and predicates on these arrows characterize the behaviors of the environment while succinct automata stay at these locations.

We now briefly explain SA_1. At location q_0, the environment transition is specified by $\varphi_1 = (\mathtt{flag}_1' = \mathtt{flag}_1 \wedge \mathtt{critical}_1' = \mathtt{critical}_1) \wedge (\mathtt{critical}_2' = 1 \Rightarrow \mathtt{flag}_2' = 1)$, meaning that SA_2 never modifies the values of $\mathtt{flag}_1$ and $\mathtt{critical}_1$ and that if SA_2 enters the critical section after the transition, denoted by $\mathtt{critical}_2' = 1$, we have $\mathtt{flag}_2' = 1$. Then, SA_1 takes a silent action to set $\mathtt{flag}_1$ to 1, meaning that it wants to enter the critical section, and enters q_1. At location q_1, the environment transition is specified by $\varphi_2 = (\mathtt{flag}_1' = \mathtt{flag}_1 \wedge \mathtt{critical}_1' = \mathtt{critical}_1) \wedge (\mathtt{critical}_2' = 1 \Rightarrow (\mathtt{flag}_2' = 1 \wedge \mathtt{turn}' = 2))$. Compared with φ_1, we see that if SA_2 enters the critical section when SA_1 is at q_1, $\mathtt{flag}_2'$ (resp. $\mathtt{turn}'$) must be 1 (resp. 2). This is because SA_2 must wait until its turn, denoted by $\mathtt{turn} = 2$, to enter the critical section once SA_1 has set $\mathtt{flag}_1$ to 1. After taking another silent action,

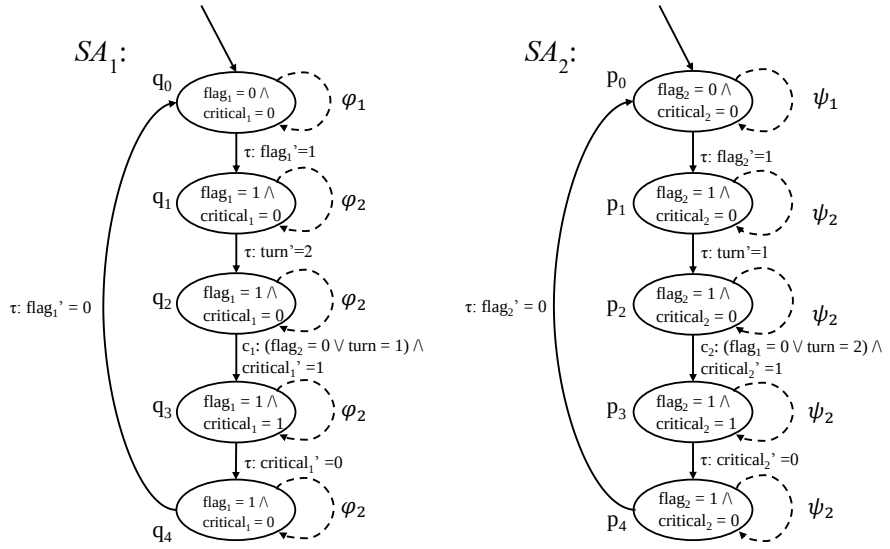

Fig. 2. Succinct automata for the Simplified Peterson's algorithm

SA_1 arrives at q_2. At location q_2, if $\texttt{flag}_2 = 0 \vee \texttt{turn} = 1$, SA_1 takes the action c_1 and enters the critical section. By entering q_4, SA_1 leaves the critical section. Finally, SA_1 resets $\texttt{flag}_1$ to 0 and comes back to q_0.

The environment transitions of SA_2 are defined by $\psi_1 = (\texttt{flag}_2' = \texttt{flag}_2 \wedge \texttt{critical}_2' = \texttt{critical}_2) \wedge (\texttt{critical}_1' = 1 \Rightarrow \texttt{flag}_1' = 1)$ and $\psi_2 = (\texttt{flag}_2' = \texttt{flag}_2 \wedge \texttt{critical}_2' = \texttt{critical}_2) \wedge (\texttt{critical}_1' = 1 \Rightarrow (\texttt{flag}_1' = 1 \wedge \texttt{turn}' = 1))$.

3.2 Parallel Composition

In rely-guarantee reasoning, the guarantee of one thread should imply the rely conditions of other threads. Similarly, we impose a compatibility condition on succinct automata running in parallel. Let q_1 (resp. q_2) be an arbitrary location in SA_1 (resp. SA_2). Informally, the compatibility condition guarantees that if SA_1 (resp. SA_2) can make a component transition from q_1 (resp. q_2) to q_1' (resp. q_2'), SA_2 (resp. SA_1) is able to mimic this transition by its environment transitions at location q_2 (resp. q_1). We formalize the *compatibility* condition as follows.

Definition 2. *SA_1 and SA_2 are compatible iff for all $(q_1, \boldsymbol{v}_G \oplus \boldsymbol{v}_{L_1}) \in S_{SA_1}, (q_2, \boldsymbol{v}_G \oplus \boldsymbol{v}_{L_2}) \in S_{SA_2}$ such that $Inv_1(q_1)(\boldsymbol{v}_G \oplus \boldsymbol{v}_{L_1})$ and $Inv_2(q_2)(\boldsymbol{v}_G \oplus \boldsymbol{v}_{L_2})$, we have*

1. *If $(q_1, \boldsymbol{v}_G \oplus \boldsymbol{v}_{L_1}) \xrightarrow{\sigma_1} (q_1', \boldsymbol{v}_G' \oplus \boldsymbol{v}_{L_1}')$, then $(q_2, \boldsymbol{v}_G \oplus \boldsymbol{v}_{L_2}) \overset{Env_2(q_2)}{\rightarrow} (q_2, \boldsymbol{v}_G' \oplus \boldsymbol{v}_{L_2})$.*
2. *If $(q_2, \boldsymbol{v}_G \oplus \boldsymbol{v}_{L_2}) \xrightarrow{\sigma_2} (q_2', \boldsymbol{v}_G' \oplus \boldsymbol{v}_{L_2}')$, then $(q_1, \boldsymbol{v}_G \oplus \boldsymbol{v}_{L_1}) \overset{Env_1(q_1)}{\rightarrow} (q_1, \boldsymbol{v}_G' \oplus \boldsymbol{v}_{L_1})$.*

Succinct automata running in parallel execute their component transitions in an interleaved manner. The formal definition of parallel composition of compatible succinct automata is defined as follows.

Definition 3. *Let* $SA_1 = (Q_1, q_0^1, V_G \cup V_{L_1}, Init_1, Inv_1, Env_1, \Sigma_1, Edge_1)$ *and* $SA_2 = (Q_2, q_0^2, V_G \cup V_{L_2}, Init_2, Inv_2, Env_2, \Sigma_2, Edge_2)$ *be two compatible succinct automata. The parallel composition of* SA_1 *and* SA_2 *is a succinct automaton* $SA_1 \parallel SA_2 = (Q, q_0, V_G \cup V_L, Init, Inv, Env, \Sigma, Edge)$*, where*

- $Q = Q_1 \times Q_2$, $q_0 = (q_0^1, q_0^2)$, $V_L = V_{L_1} \cup V_{L_2}$ *and* $\Sigma = \Sigma_1 \cup \Sigma_2$.
- $Init = Init_1 \wedge Init_2$.
- $Inv((q_1, q_2)) = Inv_1(q_1) \wedge Inv_2(q_2)$ *for each* $q_1 \in Q_1$ *and* $q_2 \in Q_2$.
- $Env((q_1, q_2)) = Env_1(q_1) \wedge Env_2(q_2)$ *for each* $q_1 \in Q_1$ *and* $q_2 \in Q_2$.
- $((q_1, q_2), \sigma, \mu, (q_1', q_2')) \in Edge$ *iff either:*
  1. *there exists an edge* $(q_1, \sigma, \mu, q_1') \in Edge_1$ *and* $q_2 = q_2'$*, or*
  2. *there exists an edge* $(q_2, \sigma, \mu, q_2') \in Edge_2$ *and* $q_1 = q_1'$.

After parallel composition, SA_1 and SA_2 share a common environment. The environment of $SA_1 \parallel SA_2$ for location (q_1, q_2) is the conjunction of the environments of SA_1 and SA_2 for location q_1 and q_2 respectively. Also, $SA_1 \parallel SA_2$ can make a component transition from (q_1, q_2) to (q_1', q_2') iff either SA_1 can make a component transition from q_1 to q_1' or the other way around. Intuitively, for each finite run of the parallel composition of two compatible succinct automata, there is a corresponding finite run in each of its components.

4 Compositional Reasoning for Succinct Automata

4.1 Safety Verification of Succinct Automata

Safety properties require that bad things should not happen. Invariants are a particular kind of safety properties that are useful in specifications. For example, the mutual exclusion property is an invariant which specifies that no more than one thread is in its critical section at any time. We introduce compositional reasoning methods for invariant verification of succinct automata and checking other safety properties can be reduced to invariant verification.

Recall that a predicate $\lambda \in \mathcal{F}(V)$ is an invariant of a transition system T if λ is satisfied on all reachable states of T. Unlike in a transition system, we have two kinds of transitions, local and environment. The way we treat them leads us to define two types of invariants of succinct automata, strong and weak. When treating both kinds of transitions equally, we reach the notion of strong invariants.

Definition 4. *A predicate* $\lambda \in \mathcal{F}(V)$ *is a strong invariant of* SA *if* λ *is satisfied on all reachable states of* SA*.*

When focusing on runs of succinct automata where environment transitions preserve λ, we reach the notion of weak invariants. The intention of weak invariants is as follows: For a program T modelled as SA, if λ is a weak invariant of SA, then, running in any environment that preserves λ, T can guarantee that λ is preserved in all its local transitions.

Definition 5. *A predicate* $\lambda \in \mathcal{F}(V)$ *is a weak invariant of* SA *if* λ *is satisfied on all runs of* SA *where environment transitions preserve* λ*.*

The notion of weak invariants is more general than strong invariants. In the following, we focus on compositionality of weak invariants. We first impose a noninterference condition on local weak invariants. This condition is to guarantee that local transitions of any component that preserve its own local weak invariant cannot invalidate local weak invariants of other components. Let λ_1 (resp. λ_2) be a weak invariant of SA_1 (resp. SA_2). Formally, we use $noninterfere(\lambda_1, \lambda_2)$ to mean the following condition: $((\lambda_1 \wedge \lambda_2 \wedge \lambda_1') \Rightarrow \lambda_2[V_G'/V_G]) \wedge ((\lambda_1 \wedge \lambda_2 \wedge \lambda_2') \Rightarrow \lambda_1[V_G'/V_G])$, where λ_i' is derived from λ_i by substituting all its variables with corresponding primed variables and $\lambda_i[V_G'/V_G]$ is derived by substituting all global variables in V_G with corresponding primed variables in V_G' for $i = 1, 2$. The parallel compositionality of weak invariants of succinct automata are formalized in the following theorem, which says that local weak invariants satisfied by all the components of the parallel composition of succinct automata guarantee a global weak invariant satisfied by the entire system as long as local weak invariants satisfy the noninterference condition.

Theorem 1. *Let SA_1 and SA_2 be compatible. Assume that $noninterfere(\lambda_1, \lambda_2)$ and λ_1 (resp. λ_2) is a weak invariant of SA_1 (resp. SA_2). We have that $\lambda_1 \wedge \lambda_2$ is a weak invariant of $SA_1||SA_2$.*

Example 2 To show that the simplified Peterson's algorithm in Figure 1 guarantees mutual exclusion, we check whether $\texttt{critical}_1 = 0 \vee \texttt{critical}_2 = 0$ is a weak invariant of $SA_1||SA_2$ in Figure 2. We define λ_1 and λ_2 by $\lambda_1 = \lambda_2 = (\texttt{critical}_1 = 0 \vee \texttt{critical}_2 = 0)$. We have verified that λ_1 (resp. λ_2) is a weak invariant of SA_1 (resp. SA_2). Also, it is easy to see that $noninterfere(\lambda_1, \lambda_2)$ holds trivially as $\lambda_1 = \lambda_2$. According to Theorem 1, we know that $\texttt{critical}_1 = 0 \vee \texttt{critical}_2 = 0$ is a weak invariant of $SA_1||SA_2$, which implies that P_1 and P_2 in Figure 1 cannot be in the critical section at the same time.

Example 3 We show the correctness of the abstract concurrent GCD programs (T_1 and T_2) in Figure 3(a). (The code is taken from [30].) To check that $T_1||T_2$ really compute the greatest common devisor (gcd) of variables a and b, we first model T_1 (resp. T_2) as SA_1 (resp. SA_2). The construction of SA_1 is shown in Figure 4 (left), where $\varphi = (a' = a) \wedge (a < b \vee b' = b)$. We omit the construction of SA_2 due to space limitation.

For convenience, we introduce two auxiliary variables A and B to SA_1 and SA_2. The value of A (resp. B) equals to the initial value of the input variable a (resp. b) and remain unchanged. Let $\lambda_1 = \lambda_2 = (\mathbf{gcd}(a, b) = \mathbf{gcd}(A, B))$, where **gcd** is a function that returns the gcd of its input. We have verified that λ_1 (resp. λ_2) is a weak invariant of SA_1 (resp. SA_2). Also, it is easy to see that $noninterfere(\lambda_1, \lambda_2)$ holds. According to Theorem 1, we know that $\mathbf{gcd}(a, b) = \mathbf{gcd}(A, B)$ is a weak invariant of $SA_1||SA_2$, which implies that $T_1||T_2$ really compute the gcd of the input values of a and b.

4.2 Simulations of Succinct Automata

We define weak simulations between succinct automata as follows.

T_1:
```
m:=0;
while (m=0) {
   atomic {
     if (a=b)
       m:=1;
     if (a>b)
       a:=a-b;
   }
}
```

T_2:
```
n:=0;
while (n=0) {
   atomic {
     if (a=b)
       n:=1;
     if (a<b)
       b:=b-a;
   }
}
```

(a) Abstract GCD Programs

T_1':
```
m:=0;
while (m=0) {
   x1:=a;
   x2:=b;
   if (x1=x2)
      m:=1;
   if (x1>x2)
      a:=x1-x2;
}
```

T_2':
```
n:=0;
while (n=0) {
   y1:=a;
   y2:=b;
   if (y1=y2)
      n:=1;
   if (y1<y2)
      b:=y2-y1;
}
```

(b) Concrete GCD Programs

Fig. 3. Concurrent GCD Programs

Definition 6. *A binary relation $\theta \subseteq S_{SA_1} \times S_{SA_2}$ is a weak simulation for (SA_1, SA_2) w.r.t. a precondition $\kappa \in \mathcal{F}(V_1 \cup V_2)$ and an invariant $\iota \in \mathcal{F}(V_1 \cup V_2)$, denoted by $SA_1 \preceq_\theta^{(\kappa,\iota)} SA_2$, iff we have the following:*

1. *$\kappa(\boldsymbol{v}_1, \boldsymbol{v}_2)$ implies $\theta((q_1, \boldsymbol{v}_1), (q_2, \boldsymbol{v}_2))$, where both q_1 and q_2 are initial.*
2. *$\theta((q_1, \boldsymbol{v}_1), (q_2, \boldsymbol{v}_2))$ implies $\iota(\boldsymbol{v}_1, \boldsymbol{v}_2)$, $Inv_1(q_1)(\boldsymbol{v}_1)$, $Inv_2(q_2)(\boldsymbol{v}_2)$ and the following:*
   - **a.** *if $(q_1, \boldsymbol{v}_1) \xrightarrow{Env_1(q_1)} (q_1, \boldsymbol{v}_1')$ and $(q_2, \boldsymbol{v}_2) \xrightarrow{Env_2(q_2)^*} (q_2, \boldsymbol{v}_2')$ and $\iota(\boldsymbol{v}_1', \boldsymbol{v}_2')$, then we have that $\theta((q_1, \boldsymbol{v}_1'), (q_2, \boldsymbol{v}_2'))$.*
   - **b.** *if $(q_1, \boldsymbol{v}_1) \xrightarrow{\sigma_1} (q_1', \boldsymbol{v}_1')$ and $\sigma_1 \neq \tau$, then there exist $(q_2', \boldsymbol{v}_2') \in S_{SA_2}$ and $\sigma_2 \in \Sigma_2$ such that $\sigma_2 = \sigma_1$, $(q_2, \boldsymbol{v}_2) \xrightarrow{\tau^*\sigma_2\tau^*} (q_2', \boldsymbol{v}_2')$ and $\theta((q_1', \boldsymbol{v}_1'), (q_2', \boldsymbol{v}_2'))$.*
   - **c.** *if $(q_1, \boldsymbol{v}_1) \xrightarrow{\tau} (q_1', \boldsymbol{v}_1')$, then there exists $(q_2', \boldsymbol{v}_2') \in S_{SA_2}$ such that $(q_2, \boldsymbol{v}_2) \xrightarrow{\tau^*} (q_2', \boldsymbol{v}_2')$ and $\theta((q_1', \boldsymbol{v}_1'), (q_2', \boldsymbol{v}_2'))$.*

Conditions 2.**b** and 2.**c** constrain local behaviors of SA_1 and SA_2 and are similar to standard notions of weak simulations [32]. Condition 2.**a** constrains the environments of the two succinct automata and requires that the weak simulation should not be affected by the environments as long as the valuations of variables in V_1 and V_2 are related by ι. Note that if we were to require that an environment transition by q_1 is simulated by zero or more environment transitions by q_2, the resulting simulation relation would not be compositional under parallel composition. Our way of dealing with environments in defining simulation or bi-simulation relations is not without precedent. For example, in process calculi, e.g., higher-order calculi [37] or cryptographic calculi [2], environments are treated separately from local transitions, and one typically requires certain relations to hold between the environments, e.g., as in the relation ι we have above. Condition 2.**a** is the key for compositionality in our notion of weak simulation.

Given κ and ι, we say that SA_1 is *weakly simulated by* SA_2 (or SA_1 *weakly refines* SA_2) *with respect to* κ *and* ι, denoted by $SA_1 \preceq^{(\kappa,\iota)} SA_2$, if there exists a weak simulation θ such that $SA_1 \preceq_\theta^{(\kappa,\iota)} SA_2$. We say that SA_1 is *weakly simulated by* SA_2, denoted by $SA_1 \preceq SA_2$, if there exist κ and ι such that $SA_1 \preceq^{(\kappa,\iota)} SA_2$. The relation $\preceq$ on succinct automata is reflexive but not transitive. However, the relation $\preceq$ on closed succinct automata is transitive. This allows us to chain together two refinement steps when reasoning about simulations between closed succinct automata.

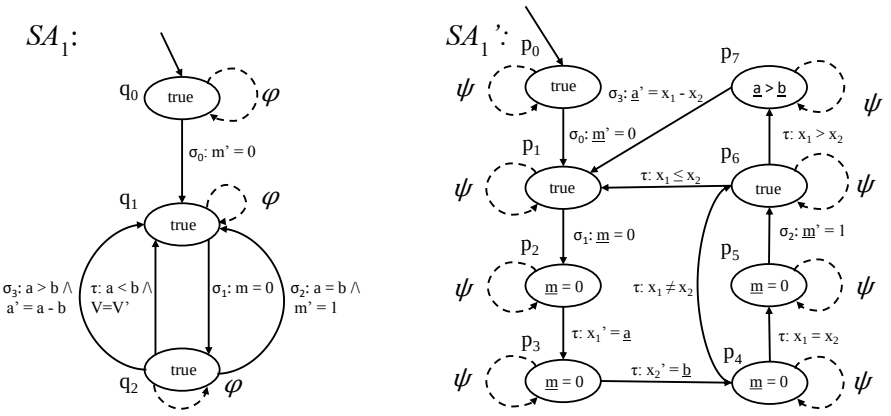

Fig. 4. Succinct automata for Concurrent GCD

Theorem 2. *The relation $\preceq$ on closed succinct automata forms a pre-order.*

For succinct automata that are not closed, we can still chain together successive refinement steps if the environment transitions of related succinct automata satisfy certain condition. We formalize this in the following theorem.

Theorem 3. *Assume that $SA_1 \preceq_{\theta_1}^{(\kappa_1,\iota_1)} SA_2$ and $SA_2 \preceq_{\theta_2}^{(\kappa_2,\iota_2)} SA_3$. Let $\kappa, \iota \in \mathcal{F}(V_1 \cup V_3)$ be predicates such that $\kappa(\boldsymbol{v}_1, \boldsymbol{v}_3)$ (resp. $\iota(\boldsymbol{v}_1, \boldsymbol{v}_3)$) iff there exists $\boldsymbol{v}_2$ such that $\kappa_1(\boldsymbol{v}_1, \boldsymbol{v}_2) \wedge \kappa_2(\boldsymbol{v}_2, \boldsymbol{v}_3)$ (resp. $\iota_1(\boldsymbol{v}_1, \boldsymbol{v}_2) \wedge \iota_2(\boldsymbol{v}_2, \boldsymbol{v}_3)$). We have that $SA_1 \preceq_{\theta_2 \circ \theta_1}^{(\kappa,\iota)} SA_3$ if the following holds: Assume that $(q_1, \boldsymbol{v}_1) \overset{Env_1(q_1)}{\longrightarrow} (q_1, \boldsymbol{v}_1')$, $(q_3, \boldsymbol{v}_3) \overset{Env_3(q_3)^*}{\longrightarrow} (q_3, \boldsymbol{v}_3')$, $\iota(\boldsymbol{v}_1, \boldsymbol{v}_3)$ and $\iota(\boldsymbol{v}_1', \boldsymbol{v}_3')$. For any $\boldsymbol{v}_2$ such that $\iota_1(\boldsymbol{v}_1, \boldsymbol{v}_2) \wedge \iota_2(\boldsymbol{v}_2, \boldsymbol{v}_3)$ and for all $q_2 \in Q_2$, there exists $\boldsymbol{v}_2'$ such that $(q_2, \boldsymbol{v}_2) \overset{Env_2(q_2)}{\longrightarrow} (q_2, \boldsymbol{v}_2')$ and $\iota_1(\boldsymbol{v}_1', \boldsymbol{v}_2') \wedge \iota_2(\boldsymbol{v}_2', \boldsymbol{v}_3')$.*

Given $\theta_1 \subseteq S_{SA_1} \times S_{\widehat{SA_1}}$ and $\theta_2 \subseteq S_{SA_2} \times S_{\widehat{SA_2}}$, we define $\theta_1 \otimes \theta_2 \subseteq S_{SA_1||SA_2} \times S_{\widehat{SA_1}||\widehat{SA_2}}$ as

$$\theta_1 \otimes \theta_2(((q_1,q_2),\mathbf{v}),((\widehat{q_1},\widehat{q_2}),\widehat{\mathbf{v}})) \texttt{ iff } \theta_1((q_1,\mathbf{v}_G \oplus \mathbf{v}_{L_1}),(\widehat{q_1},\widehat{\mathbf{v}_G} \oplus \widehat{\mathbf{v}_{L_1}})) \texttt{ and } \theta_2((q_2,\mathbf{v}_G \oplus \mathbf{v}_{L_2}),(\widehat{q_2},\widehat{\mathbf{v}_G} \oplus \widehat{\mathbf{v}_{L_2}})),$$

where $\mathbf{v} = \mathbf{v}_G \oplus \mathbf{v}_{L_1} \oplus \mathbf{v}_{L_2}$ and $\widehat{\mathbf{v}} = \widehat{\mathbf{v}_G} \oplus \widehat{\mathbf{v}_{L_1}} \oplus \widehat{\mathbf{v}_{L_2}}$.

To ensure compositionality of weak simulations, we also impose a noninterference condition on ι_1 and ι_2 here. We reuse $noninterfere(\iota_1, \iota_2)$ to denote the following condition: $((\iota_1 \wedge \iota_2 \wedge \iota_1') \Rightarrow \iota_2[V_G'/V_G][\widehat{V_G}'/\widehat{V_G}]) \wedge ((\iota_1 \wedge \iota_2 \wedge \iota_2') \Rightarrow \iota_1[V_G'/V_G][\widehat{V_G}'/\widehat{V_G}])$. The following theorem shows that weak simulations of succinct automata are preserved under parallel composition.

Theorem 4. *Assume that SA_1 (resp. $\widehat{SA_1}$) and SA_2 (resp. $\widehat{SA_2}$) are compatible and that $noninterfere(\iota_1, \iota_2)$. We have that $SA_1 \preceq_{\theta_1}^{(\kappa_1,\iota_1)} \widehat{SA_1}$ and $SA_2 \preceq_{\theta_2}^{(\kappa_2,\iota_2)} \widehat{SA_2}$ implies $SA_1||SA_2 \preceq_{\theta_1 \otimes \theta_2}^{(\kappa_1 \wedge \kappa_2, \iota_1 \wedge \iota_2)} \widehat{SA_1}||\widehat{SA_2}$.*

Example 4 We show that the abstract concurrent GCD programs (T_1 and T_2) in Figure 3(a) are refined by the concrete GCD programs (T_1' and T_2') in Figure 3(b). The bodies of the while loops in T_1 and T_2 are executed atomically and are refined to corresponding code in T_1' and T_2' to allow interleaving.

In Figure 4, we model thread T_1 (resp. T_1') as SA_1 (resp. SA_1'), where $\varphi = (a' = a) \wedge (a < b \vee b' = b)$ and $\psi = (\underline{a}' = \underline{a}) \wedge (\underline{a} < \underline{b} \vee \underline{b}' = \underline{b})$. Let κ_1 and ι_1 be defined by $\kappa_1 = (\underline{a} = a \wedge \underline{b} = b \wedge \underline{m} = m)$ and $\iota_1 = (\underline{a} = a \wedge \underline{b} = b \wedge \underline{m} = m)$. In our experiment, using the verification tool we have implemented, we have verified that $SA_1' \preceq_{\theta_1}^{(\kappa_1,\iota_1)} SA_1$ holds for some θ_1. Similarly, we have modeled T_2 (resp. T_2') as SA_2 (resp. SA_2') and checked in our experiment that $SA_2 \preceq_{\theta_2}^{(\kappa_2,\iota_2)} SA_2'$ holds for some θ_2. By Theorem 4, we have that $SA_1||SA_2 \preceq_{\theta_1\otimes\theta_2}^{(\kappa_1\wedge\kappa_2,\iota_1\wedge\iota_2)} SA_1'||SA_2'$.

4.3 Safety Property Preservation under Refinement

It is obvious that strong invariants are preserved under refinements. We show in the following that weak invariants of succinct automata are also preserved under refinements.

We write $\mathcal{WS}(\theta, Env_1, Env_2)$ to mean that: if $\theta((q_1, \mathbf{v}_1), (q_2, \mathbf{v}_2))$ and $(q_1, \mathbf{v}_1) \xrightarrow{Env_1(q_1)} (q_1, \mathbf{v}_1')$, there exists $(q_2, \mathbf{v}_2') \in S_{SA_2}$ such that $(q_2, \mathbf{v}_2) \xrightarrow{Env_2(q_2)^*} (q_2, \mathbf{v}_2')$ and $\theta((q_1, \mathbf{v}_1'), (q_2, \mathbf{v}_2'))$. If $\mathcal{WS}(\theta, Env_1, Env_2)$ holds, for each run in SA_1, we can construct a corresponding run in SA_2 such that the two runs are related by θ. Thus, we have the following lemma that links reachability and weak simulations.

Lemma 1. *Assume that $SA_1 \preceq_\theta^{(\kappa,\iota)} SA_2$, where $\mathcal{WS}(\theta, Env_1, Env_2)$ holds. For all states $(q_1, \boldsymbol{v}_1) \in S_{SA_1}$ and $(q_2, \boldsymbol{v}_2) \in S_{SA_2}$ such that $\theta((q_1, \boldsymbol{v}_1), (q_2, \boldsymbol{v}_2))$, if $(q_1, \boldsymbol{v}_1) \rightarrow^* (q_1', \boldsymbol{v}_1')$ for some $(q_1', \boldsymbol{v}_1') \in S_{SA_1}$, there exists $(q_2', \boldsymbol{v}_2') \in S_{SA_2}$ such that $(q_2, \boldsymbol{v}_2) \rightarrow^* (q_2', \boldsymbol{v}_2')$ and $\theta((q_1', \boldsymbol{v}_1'), (q_2', \boldsymbol{v}_2'))$.*

As invariants verification can be reduced to reachability problems, we can prove by contradiction that the following theorem holds.

Theorem 5. *Assume that $SA_1 \preceq_\theta^{(\kappa,\iota)} SA_2$, where $\mathcal{WS}(\theta, Env_1, Env_2)$ holds, and for each initial state $(q_1, \boldsymbol{v}_1) \in S_{SA_1}$, there exists an initial state $(q_2, \boldsymbol{v}_2) \in S_{SA_2}$ such that $\kappa(\boldsymbol{v}_1, \boldsymbol{v}_2)$. Let $\lambda_1 \in \mathcal{F}(V_1)$ and $\lambda_2 \in \mathcal{F}(V_2)$ be two predicates such that $\neg\lambda_1(\boldsymbol{v}_1) \wedge \iota(\boldsymbol{v}_1, \boldsymbol{v}_2)$ implies $\neg\lambda_2(\boldsymbol{v}_2)$. If λ_2 is a weak invariant of SA_2, then λ_1 is a weak invariant of SA_1.*

Example 5 We give a short example to show that $\mathbf{gcd}(a, b) = \mathbf{gcd}(A, B)$ is a weak invariant of the concrete GCD programs, which implies that the concrete GCD programs also compute the gcd of the input variables. First, we know from Example 3 that $\mathbf{gcd}(a, b) = \mathbf{gcd}(A, B)$ is a weak invariant of the abstract GCD programs. Second, we know from Example 4 that the concrete GCD programs refine the abstract GCD programs. Hence, from Theorem 5, we can prove that $\mathbf{gcd}(a, b) = \mathbf{gcd}(A, B)$ is also a weak invariant of the concrete GCD programs.

5 Automatic Verification of Succinct Automata

5.1 Automatic Verification of Infinite State Succinct Automata

We propose an SMT-based approach (Algorithm 1) to checking weak simulations between infinite state succinct automata. Our approach is developed under the assumption that environment transitions at each location of the abstract succinct automata are transitive. In the rest of this section, we assume that $Env_2(q_2)$ is transitive for all $q_2 \in Q_2$ and SA_1 and SA_2 use different set of variables.

Proving $SA_1 \preceq^{(\kappa,\iota)} SA_2$ amounts to showing the existence of a simulation relation θ such that $SA_1 \preceq_\theta^{(\kappa,\iota)} SA_2$. We define first order formulas $\Psi_{(q_1,q_2)}$ over $V_1 \cup V_2$ for a set of pairs of locations $(q_1, q_2) \in Q_1 \times Q_2$. The intention is that when our algorithm terminates, we can construct a relation θ by defining $\theta((q_1, \mathbf{v}_1), (q_2, \mathbf{v}_2))$ iff $\Psi_{(q_1,q_2)}(\mathbf{v}_1, \mathbf{v}_2)$ such that θ satisfies Condition 2 in Definition 6.

Our algorithm follows the basic fixed point iteration method [14]. The main function in Algorithm 1 is $Refine\text{-}Weak\text{-}Simulation$. It first computes a set Θ that contains all the pairs (q_1, q_2) for which we need to define constraints. Then, it defines the initial value of $\Psi_{(q_1,q_2)}$ for each $(q_1, q_2) \in \Theta$. In each fixed point iteration (Line 38-41), we first generate constraints by calling function $GenConstraints$ and then update the value of $\Psi_{(q_1,q_2)}$ through function $UpdatePsi$ according to the satisfiability of the constraints generated for (q_1, q_2). When the greatest fixed point is reached, we check whether Condition 1 in Definition 6 is satisfied (Line 42).

Due to space limitation, we omit the pseudo code for the function $GenPairs$ (called in Line 34) and explain it briefly as follows. Let $\Pi_\sigma(q)$ denote the set of finite local paths π such that π starts from q and the action labels along π are of the form $\tau^*\sigma\tau^*$ (resp. τ^*), when $\sigma \neq \tau$ (resp. $\sigma = \tau$). $GenPairs$ is a recursive function which takes a set Θ of pairs of locations as input and returns another set of pairs of locations. Let Θ' be an empty set. First, for each $(q_1, q_2) \in \Theta$, it adds to Θ' the set of (q_1', q_2') such that there exists an edge (q_1, σ, μ, q_1') and a path $\pi \in \Pi_\sigma(q_2)$ that ends in q_2'. Then, $GenPairs$ makes a recursive call $GenPairs(\Theta' \backslash \Theta)$ and returns $\Theta \cup GenPairs(\Theta' \backslash \Theta)$.

Function $GenConstraints$ generates three constraints $\neg C_1, \neg C_2$ and $\neg C_3$ for each $(q_1, q_2) \in \Theta$. Here, we use formulas C_1 and C_2 (Line 3-4) to specify Condition 2.**a** in Definition 6. Since we assume that $Env_2(q_2)$ is transitive, C_1 and C_2 can characterize all the cases in Condition 2.**a**. Formula C_3 (Line 9) are used to specify Condition 2.**b** and 2.**c**. We use $WP(e, \Psi)$ (resp. $WP(\pi, \Psi)$) to denote the weakest precondition such that Ψ holds after taking a component transition (resp. a sequence of component transitions) by executing e (resp. π).

Function $UpdatePsi$ checks the satisfiability of all the constraints generated by $GenConstraints$. If a constraint $\neg(\Psi_{(q_1,q_2)} \Rightarrow \Phi)$ is satisfiable, $\Psi_{(q_1,q_2)}$ fails to satisfy Condition 2 in Definition 6. In this case, we strengthen $\Psi_{(q_1,q_2)}$ in Line 21, 23 or 25 depending on the type of the constraint. Here, type i constraints refers to those of the form $\neg C_i$ generated by $GenConstraints$ for $i = 1, 2, 3$.

5.2 Experimental Results

We have prototyped our SMT-based algorithm in F# and used Z3 [15] in our implementation. We applied our tool to check multiple weak simulations between concurrent

Algorithm 1: An Algorithm to Check Weak Simulations of Succinct Automata

Input: SA_1 and SA_2 and parameters κ and ι.

Output: If the algorithm return *Yes*, $SA_1 \preceq^{(\kappa,\iota)} SA_2$ holds. If the algorithm returns *No*, $SA_1 \preceq^{(\kappa,\iota)} SA_2$ does not hold.

Function $GenConstraints\,(SA_1, SA_2, \Theta, \iota)$ **is**
  foreach $(q_1, q_2) \in \Theta$ **do**
    $C_1 := \Psi_{(q_1,q_2)} \Rightarrow ((\Phi_{Env_1(q_1)} \wedge \iota[V_1'/V_1]) \Rightarrow \Psi_{(q_1,q_2)}[V_1'/V_1])$
    $C_2 := \Psi_{(q_1,q_2)} \Rightarrow ((\Phi_{Env_1(q_1)} \wedge \Phi_{Env_2(q_2)} \wedge \iota[V_1'/V_1][V_2'/V_2]) \Rightarrow$
      $\Psi_{(q_1,q_2)}[V_1'/V_1][V_2'/V_2])$
      where $\Phi_{Env_i(q_i)} = Inv_i(q_i) \wedge Env_i(q_i) \wedge Inv_i'(q_i) \wedge V_{L_i} = V_{L_i}'$
    $constraints := constraints \cup \{\neg C_1, \neg C_2\}$
    foreach $e = (q_1, \sigma, \mu, q_1') \in Edge_1$ **do**
      $C_3 := \Psi_{(q_1,q_2)} \Rightarrow (G \Rightarrow (WP(e', \bigvee_{\pi \in \Pi_\sigma(q_2)} WP(\pi, \Psi_{(q_1',q_2')}))))$
        where G is the guard of e, e' is derived from e by substituting its guard
        with $True$, and π ends at location q_2'
      $constraints := constraints \cup \{\neg C_3\}$
  return $constraints$

Function $UpdatePsi\,(constraints, V_1, V_2, \Theta)$ **is**
  foreach $(q_1, q_2) \in \Theta$ **do**
    $\Psi'_{(q_1,q_2)} := \Psi_{(q_1,q_2)}$
  foreach $\neg(\Psi_{(q_1,q_2)} \Rightarrow \Phi) \in constraints$ **do**
    if $\neg(\Psi_{(q_1,q_2)} \Rightarrow \Phi)$ *is satisfiable* **then**
      if $\neg(\Psi_{(q_1,q_2)} \Rightarrow \Phi)$ *is a type* 1 *constraint* **then**
        $\Psi'_{(q_1,q_2)} := \Psi'_{(q_1,q_2)} \wedge \forall V_1'\ \Phi$
      if $\neg(\Psi_{(q_1,q_2)} \Rightarrow \Phi)$ *is a type* 2 *constraint* **then**
        $\Psi'_{(q_1,q_2)} := \Psi'_{(q_1,q_2)} \wedge \forall V_1', V_2'\ \Phi$
      if $\neg(\Psi_{(q_1,q_2)} \Rightarrow \Phi)$ *is a type* 3 *constraint* **then**
        $\Psi'_{(q_1,q_2)} := \Psi'_{(q_1,q_2)} \wedge \Phi$
  foreach $(q_1, q_2) \in \Theta$ **do**
    $\Psi_{(q_1,q_2)} := \Psi'_{(q_1,q_2)}$
  if *none of the constraints are satisfiable* **then**
    return *Fixed Point Reached*
  else
    return *Continue Iteration*

Function `Refine-Weak-Simulation` $(SA_1, SA_2, \kappa, \iota)$ **is**
  $\Theta := GenPairs(\{(q_{init_1}, q_{init_2})\})$
  foreach $(q_1, q_2) \in \Theta$ **do**
    $\Psi_{(q_1,q_2)} := \iota \wedge Inv_1(q_1) \wedge Inv_2(q_2)$
  $constraints := \emptyset$
  repeat
    $constraint := GenConstraints(SA_1, SA_2, \Theta, \iota)$
    $result := UpdatePsi(constraints, V_1, V_2, \Theta)$
  until $result =$ *Fixed Point Reached*
  if $\kappa \Rightarrow \Psi_{(q_{init1}, q_{init2})}$ *is valid* **then**
    return *Yes*
  else
    return *No*

Experiments\Description	Number of Iterations	Time: sec	Verification Results
1. Concurrent GCD	8, 8	3.299, 2.709	Yes, Yes
2. Peterson's Algorithm	5	0.937	No
3. Invariant Hoisting	2	0.149	Yes
4. Strength Reduction (SR)	7	0.628	Yes
5. Induction Variable Elimination (IVE)	3	0.227	Yes
6. SR combined with IVE	3	0.188	Yes

Table 1. Experimental Results for Refinement Checking between Concurrent Programs

Source Program C_1

```
while(i<n) {
   t:=x+1;
   i:=i+t;
}
```

Target Program C_2

```
t:=x+1;
while(i<n) {
   i:=i+t;
}
```

Fig. 5. Invariant Hoisting

programs. The experimental results are summarized in Table 1, where the second column lists the number of iterations to reach the greatest fixed point in each experiment, the third column refers to the entire time it takes to finish the refinement checking and the last column indicates whether we found the simulation relations in the experiment.

In the first experiment, we focused on the refinement problem in Example 4 and verified that the concrete GCD programs in Figure 3(b) do refine the abstract concurrent GCD programs in Figure 3(a). As there are two threads in GCD programs, the experimental results in the second row are listed in pairs. The first (resp. second) element of the pairs are related to checking refinements of the first (resp. second) threads.

In the second experiment, we switched the order of the first two instructions in the while loop of the simplified Peterson's algorithm in Figure 1 and have found that the modified loop body does not refine the original one. Actually, the modified Peterson's algorithm does not guarantee mutual exclusion. This is consistent with our theorems because if the modified Peterson's algorithm refines the one in Figure 1, it should also guarantee mutual exclusion. The experimental results in the third row are about checking whether the modified loop body of P_1 in Figure 1 refines the original one.

In the rest of the experiments, we checked the correctness of several compiler optimizations on concurrent programs. The compiler optimizations involved are invariant hoisting, strength reduction and induction variable elimination, which are often used for sequential programs in practice [1]. In concurrent settings, due to unintended behaviors of the environments, these optimizations are not always correct. Actually, they are correct only when the environments are reasonably constrained. In our experiment, we made valid assumptions about the environments of the verified programs and modelled the environments using environment transitions of succinct automata. The setting of Experiments 3-6 is summarized as follows. In each experiment, a source program C is transformed to a target program C' by compiler optimizations, and the assumptions on the environments of C and C' are made. We model C (resp. C') as SA (resp. SA'). We say that C is correctly transformed to C' (under the given assumptions about the environments) if $SA' \preceq^{(\kappa,\iota)} SA$ holds, where both the precondition κ and the invariant ι specify that common variables in SA and SA' should have the same values. All the example code in these experiments are from [30].

Source Program C

```
i:=0;
while(i<n) {
   x:=x+6*i;
   i:=i+1;
}
```

Middle Program C_1

```
i:=0;
k:=0;
while(i<n) {
   x:=x+k;
   i:=i+1;
   k:=k+6;
}
```

Target Program C_2

```
k:=0;
r:=6*n;
while(k<r) {
   x:=x+k;
   k:=k+6;
}
```

Fig. 6. Strength Reduction and Induction Variable Elimination

Experiment 3 concerns invariant hoisting. By invariant hoisting, compilers hoist loop invariants out of a loop to improve runtime performance since the hoisted loop invariants do not have to be executed in each iteration of the loop. The source program C_1 in Figure 5 is transformed to the target program C_2 in Figure 5 by hoisting the loop invariant `t:=x+1`. Both C_1 and C_2 are supposed to run in an environment that does not modify the value of x and t. We have verified that C_2 correctly refines C_1 in the assumed environments.

Experiments 4 and 5 concern strength reduction and induction variable elimination. Through strength reduction, compilers improve runtime performance by replacing expensive computations with equivalent but less expensive ones. By induction variable elimination, compilers reduce the number of computations in a loop. First, the source program C in Figure 6 is transformed to a middle program C_1 in Figure 6 by strength reduction, where a local variable k is introduced and multiplication is replaced by addition. We assume that the environment of C does not modify i and the environment of C_1 does not modify i and k. We verified the correctness of this transformation in Experiment 4. Then, C_1 is transformed to a target program C_2 in Figure 6 by induction variable elimination, where the induction variable i is eliminated and the variable k is used as the new induction variable. We assume that the environment of C_2 does not modify n, r and k. We verified the correctness of this refinement in Experiment 5.

Instead of using Theorem 3 to prove that C in Figure 6 is correctly transformed to C_2 in Figure 6, we went one step further in Experiment 6 to see whether our tool can directly verify the refinement between C and C_2 in Figure 6. As we can see in Table 1, it even takes less time to check the refinement between C and C_2 directly. The main reason for this is that C_1 is a bit more complex than both C and C_2. Therefore, checking the refinement between C and C_1 and that between C_1 and C_2 are both more expensive.

6 Conclusions and Future Work

We have introduced a framework for formal verifications of shared-variable concurrent programs. In our framework, safety verification and simulations of concurrent programs are parallel compositional and algorithmic. Succinct automata-based approaches can be applied to extend safety verification of concurrent programs from the source code level down to the binary level in a compositional way. In our future work, we will focus on developing efficient algorithms for various formal verification problems, especially checking weak simulations, of succinct automata and verify refinements between larger concurrent programs.

Reference

1. Compiler optimizations. http://www.compileroptimizations.com/.
2. M. Abadi and A. D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Inf. Comput.*, 148(1):1–70, 1999.
3. C. Baier and J. Katoen. *Principles of model checking*. MIT Press, 2008.
4. S. Bensalem, A. Bouajjani, C. Loiseaux, and J. Sifakis. Property preserving simulations. In *Proceedings of the 4th International Conference on Computer Aided Verification (CAV 1992), Montreal, Canada*, pages 260–273, 1992.
5. S. D. Brookes. Full abstraction for a shared variable parallel language. In *Proceedings of the 8th Annual Symposium on Logic in Computer Science (LICS 1993), Montreal, Canada*, pages 98–109, 1993.
6. J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic model checking: 10ˆ20 states and beyond. *Inf. Comput.*, 98(2):142–170, 1992.
7. S. Burckhardt, M. Musuvathi, and V. Singh. Verifying local transformations on relaxed memory models. In *Proceedings of the 19th International Conference on Compiler Construction (CC 2010), Paphos, Cyprus*, pages 104–123, 2010.
8. E. M. Clarke, T. Filkorn, and S. Jha. Exploiting symmetry in temporal logic model checking. In *Proceedings of the 5th International Conference on Computer Aided Verification (CAV 1993), Elounda, Greece*, pages 450–462, 1993.
9. E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement. In *Proceedings of the 12th International Conference on Computer Aided Verification (CAV 2000), Chicago, IL, USA*, pages 154–169, 2000.
10. E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM*, 50(5):752–794, 2003.
11. E. M. Clarke, O. Grumberg, and D. E. Long. Model checking and abstraction. *ACM Trans. Program. Lang. Syst.*, 16(5):1512–1542, 1994.
12. A. Cohen and K. S. Namjoshi. Local proofs for global safety properties. *Formal Methods in System Design*, 34(2):104–125, 2009.
13. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th Symposium on Principles of Programming Languages (POPL 2007), Los Angeles, California, USA*, pages 238–252, 1977.
14. B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order (2. ed.)*. Cambridge University Press, 2002.
15. L. M. de Moura and N. Bjørner. Z3: an efficient SMT solver. In *Proceedings of the 14th International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS 2008), Budapest, Hungary*, pages 337–340, 2008.
16. E. A. Emerson and A. P. Sistla. Symmetry and model checking. In *Proceedings of the 5th International Conference on Computer Aided Verification (CAV 1993), Elounda, Greece*, pages 463–478, 1993.
17. C. Flanagan, S. N. Freund, and S. Qadeer. Thread-modular verification for shared-memory programs. In *Proceedings of the 11th European Symposium on Programming (ESOP 2002), Grenoble, France*, pages 262–277, 2002.
18. C. Flanagan and S. Qadeer. Thread-modular model checking. In *Proceedings of the 10th International SPIN Workshop (SPIN 2003), Portland, OR, USA*, pages 213–224, 2003.
19. P. Godefroid. Using partial orders to improve automatic verification methods. In *Proceedings of the 2nd International Conference on Computer Aided Verification (CAV 1990), New Brunswick, NJ, USA*, pages 176–185, 1990.

20. S. Graf and H. Saïdi. Construction of abstract state graphs with PVS. In *Proceedings of the 9th International Conference on Computer Aided Verification (CAV 1997), Haifa, Israel*, pages 72–83, 1997.
21. O. Grumberg and D. E. Long. Model checking and modular verification. *ACM Trans. Program. Lang. Syst.*, 16(3):843–871, 1994.
22. A. Gupta, C. Popeea, and A. Rybalchenko. Predicate abstraction and refinement for verifying multi-threaded programs. In *Proceedings of the 38th Symposium on Principles of Programming Languages (POPL 2011), Austin, TX, USA*, pages 331–344, 2011.
23. C. Hawblitzel, E. Petrank, S. Qadeer, and S. Tasiran. Automated and modular refinement reasoning for concurrent programs. In *Proceedings of the 27th International Conference on Computer Aided Verification (CAV 2015), San Francisco, CA, USA*, pages 449–465, 2015.
24. T. A. Henzinger, R. Jhala, and R. Majumdar. Race checking by context inference. In *Proceedings of the 2004 Conference on Programming Language Design and Implementation (PLDI 2004), Washington, DC, USA*, pages 1–13, 2004.
25. T. A. Henzinger, R. Jhala, R. Majumdar, and S. Qadeer. Thread-modular abstraction refinement. In *Proceedings of the 15th International Conference on Computer Aided Verification (CAV 2003), Boulder, CO, USA*, pages 262–274, 2003.
26. C. N. Ip and D. L. Dill. Better verification through symmetry. In *Proceedings of the 11th International Conference on Computer Hardware Description Languages and their Applications (CHDL 1993), Ontario, Canada*, pages 97–111, 1993.
27. C. B. Jones. Specification and design of (parallel) programs. In *IFIP Congress*, pages 321–332, 1983.
28. C. B. Jones. Tentative steps toward a development method for interfering programs. *ACM Trans. Program. Lang. Syst.*, 5(4):596–619, 1983.
29. S. Kundu, S. Lerner, and R. Gupta. Automated refinement checking of concurrent systems. In *Proceedings of the 2007 International Conference on Computer-Aided Design (ICCAD 2007), San Jose, CA, USA*, pages 318–325, 2007.
30. H. Liang, X. Feng, and M. Fu. Rely-guarantee-based simulation for compositional verification of concurrent program transformations. *ACM Trans. Program. Lang. Syst.*, 36(1):3, 2014.
31. A. Lochbihler. Verifying a compiler for java threads. In *Proceedings of the 19th European Symposium on Programming (ESOP 2010), Paphos, Cyprus*, pages 427–447, 2010.
32. R. Milner. *Communication and concurrency*. PHI Series in computer science. Prentice Hall, 1989.
33. J. Misra and K. M. Chandy. Proofs of networks of processes. *IEEE Trans. Software Eng.*, 7(4):417–426, 1981.
34. F. Nielson, H. R. Nielson, and C. Hankin. *Principles of program analysis*. Springer, 1999.
35. D. Peled. Combining partial order reductions with on-the-fly model-checking. In *Proceedings of the 6th International Conference on Computer Aided Verification (CAV 1994), Stanford, California, USA*, pages 377–390, 1994.
36. A. Pnueli. In transition from global to modular temporal reasoning about programs. In K. R. Apt, editor, *Logics and Models of Concurrent Systems*, pages 123–144. Springer-Verlag New York, Inc., New York, NY, USA, 1985.
37. D. Sangiorgi, N. Kobayashi, and E. Sumii. Environmental bisimulations for higher-order languages. In *Proceedings of the 22nd IEEE Symposium on Logic in Computer Science (LICS 2007), Wroclaw, Poland*, pages 293–302, 2007.
38. A. Valmari. A stubborn attack on state explosion. In *Proceedings of the 2nd International Conference on Computer Aided Verification (CAV 1990), New Brunswick, NJ, USA,*, pages 156–165, 1990.
39. Q. Xu, W. P. de Roever, and J. He. The rely-guarantee method for verifying shared variable concurrent programs. *Formal Asp. Comput.*, 9(2):149–174, 1997.